arXiv:1612.06079v1 [cs.DL] 19 Dec 2016

Empirical and theoretical consequences of using the Euclidean Index for assessing individual scholarly impact

Jens Peter Andersen[a,*]

[a] *Danish Centre for Studies on Research and Research Policy. Aarhus University. Bartholins Alle 7, Dk-8000 Aarhus C. Denmark.*

Abstract

The recently proposed Euclidean index offers a novel approach to measure the citation impact of academic authors, in particular as an alternative to the h-index. We test if the index provides new, robust information, not covered by existing bibliometric indicators, discuss the measurement scale and the degree of distinction between analytical units the index offers. We find that the Euclidean index does not outperform existing indicators on these topics and that the main application of the index would be solely for ranking, which is not seen as a recommended practice.

Keywords: citation analysis, evaluative bibliometrics, euclidean index, h-index, micro-level bibliometrics, author evaluation, indicator validation, indicator stability

In a recent paper, Perry & Reny (2016) propose an indicator of the impact of individuals, designed to provide a more rigorous approach hereto than the h-index which has been a de facto standard in several assessment situations, ranging from academic hiring to research grant applications. Their proposed indicator, the Euclidean index, is interpreted as the Euclidean length of the vector composed of the total citation scores of an individual's papers, or in other words the square root of the sum of squares of an author's citations per paper.

*Corresponding author
Email address: jpa@ps.au.dk (Jens Peter Andersen)

While the authors claim this to be a new indicator, it is rather a modification of two previously proposed indicators; the R-index (Jin et al., 2007) and the e-index (Zhang, 2009), which use the same general formulation as the Euclidean index, with only minor differences (further details in the methodology). Despite this, the approach used by Perry & Reny warrants a discussion of this index in particular, and the properties of these types of indices in general.

Perry & Reny designed their index based on five axioms, which they consider crucial for an indicator of an individual's citation impact, and through two empirical tests, they find that it outperforms the h-index with regard to 1) maintaining inter-field rankings after rescaling and 2) matching a ranking of top universities created by the National Research Council (NRC) of the United States of America. The approach of using rigorous, pre-defined axioms for designing scientometric indices is commendable, and the resulting index possesses some qualities, which are obviously missing in the h-index, namely their five axioms; monotonicity, independence, depth relevance, scale invariance and directional consistency (Perry & Reny, 2016). However, we will claim that there are two essential discussions which are omitted in this paper, and a number of other papers related to the h-index; the validity and ethics of assessing individuals as well as the sufficiency of the proposed axioms. We will argue for the former in the following section, followed by a case for the latter, including macro-scale empirical analysis, in the remainder of the paper.

1. Assessment of individuals

All bibliometric research is divided into three parts, one of which is evaluative bibliometrics, descriptive bibliometrics another and referencing studies a third. Evaluative and descriptive bibliometrics may focus on citations or publications and are generally quantitative in nature. For both types of quantitative studies, methodological approaches depend on the aggregation level and thus the unit size being evaluated. Macro-scale studies focus on world- or nation-wide research outputs, meso-level look at large research units, such as univer-

sities and possibly university departments, while studies of individual authors are considered micro-scale bibliometric studies. It is essential for any type of quantitative bibliometric study to use these aggregated units of analysis just as well as citation counts are aggregated references, viewing not individual works but the collective works of larger units as the object of study. By doing so, we move from the sometimes arbitrary selections of one reference over another, due to norms, personal relations, availability etc., to a generalized, statistical view (Van Raan, 1998). Whether these sociological, personal and normative variations at the reference level are actually randomly distributed over large aggregates is debatable (Waltman et al., 2013), and certainly requires a considerable amount of care, delicacy and large numbers of publications per unit of analysis Vinkler (2007); Costas et al. (2010). As a consequence, within the field of bibliometrics, analyses of individuals, or micro-level citation analysis, have focused on more descriptive elements, such as collaborative networks or factors influencing productivity (Costas et al., 2010) or as a support tool for informed peer-review (Aksnes & Taxt, 2004). Traditionally, ranking algorithms and hard impact assessments have not been applied to micro-level aggregates, such as authors, due to the sensitivity of indicators to even fairly small outliers when regarding small samples. Since the publication of the h-index (Hirsch, 2005), there has been a growing attention on the dangers of assessment on the individual level and the potential influence on the research process (Vinkler, 2007; Van Raan, 2006) but also a large number of papers utilizing and attempting to improve the h-index (Bornmann et al., 2008; Braun et al., 2006; Moed, 2009; Glänzel & Moed, 2012). Perry & Reny are also critical towards the concept of evaluating individuals through citation counts, and the h-index in particular, yet maintain that these assessments are a matter of fact:

> [Citation indices] are regularly used to inform critical decisions about funding, promotion and tenure. With decisions of this magnitude on the line, one should approach the problem of developing a good index as systematically as possible. (Perry & Reny, 2016, 2722)

We should certainly take into account those practices that are currently part of the norms of various scientific disciplines (de Rijcke & Rushforth, 2015; Rushforth & de Rijcke, 2015; Hicks et al., 2015), and the types of indicators which are used and understood outside the sciento- and bibliometric fields, e.g. by funders, managers or scientists themselves (Leydesdorff et al., 2016). In assessing and developing new indicators, we should thus consider if they are tools for use of bibliometricians or end-users. The Euclidean index is clearly aimed at the latter, and it should therefore also be clearly understandable and interpretable by these.

2. Assessment of indicators

As mentioned above, Perry & Reny apply an axiomatic approach to indicator design, basing the justification of the Euclidean index on five, mathematical requisites for a bibliometric indicator of individual, scholarly impact. The axioms are common, mathematical requirements for good measurements, and the authors argue well for the properties of the Euclidean index and how the index performs better than the h-index at ranking economics scholars. However, the authors have not tested the performance of their indicator against basic bibliometric units, nor other existing indicators, and it is thus unclear whether this new indicator even offers new or more robust information. The only empirical test performed is on the ranking ability of the indicator versus the h-index, compared to a peer-produced ranking. As it has been recommended that ranking should not be performed on the individual level (Van Raan, 2006; Vinkler, 2007; Costas et al., 2010), this practice should be avoided, and will not be pursued further in this paper, which will instead focus on the basic properties of all bibliometric indicators; size and impact. These two factors have been identified by Leydesdorff (2009) and Bollen et al. (2009) as the main categories in which the majority of bibliometric indicators can be arranged into. The h-index and its derivatives are most clearly separated from these two factors, although not necessarily for the better (Leydesdorff, 2009; Bollen et al., 2009; Waltman &

Van Eck, 2012; Perry & Reny, 2016). The introduction of a new citation-based bibliometric indicator should provide new information or insight, in order to not be redundant. Alternatively, new indicators can seek to provide more precise or robust information on aspects of scientific outputs, thus replacing existing indicators. Both the novelty and robustness of the Euclidean index will be tested empirically and theoretically in this paper; in comparison with basic units as well as the h-, $R-$ and e-indices, which are the most closely related indicators.

Also the ordinality of the Euclidean index will be discussed, in particular with regard to distinguishing between two units of analysis. This is a problematic aspect of the index, as Perry & Reny claim the index can be used not only for ranking, but also for general statements about the underlying distributions, or at least comparisons between distributions. Given the definition of the index, and typical citation distributions, we consider this unlikely, and we will discuss this using a thought experiment and a simulation on empirical data. This is especially important, as Perry & Reny present their indicator as a standalone solution, while it has been recommended that bibliometric analyses should ideally operationalize multiple indicators (Hicks et al., 2015).

2.1. Interpreting measurements

As we claim in the introductory paragraphs, it is clear that the Euclidean index is designed with end-users (grant committees, administrators, researchers) in mind, and should therefore be easily interpretable in its own right. This can be achieved by operationalizing a unit (Hand, 2004), carrying an interpretational value either on the nominal, ordinal or scale level. Simple, well-known examples of this could be gender distinctions, US dollars and degrees Celsius, each of which have an intrinsic value. Perhaps the most interesting of these is the scale unit, degrees Celsius, as this scale is humanly defined and other scales such as Kelvin and degrees Fahrenheit are equally valid. However, by designing the scale around the freezing and boiling point of water, an additional, un-arbitrary value is added. With regard to bibliometric indicators, the same is true. Basic count variables such as number of publications (P) and citations (C) are their

own units, similar to capita counts, but with e.g. average scores, this becomes less obvious. However, as we have the basic units P and C, we can also express the unit of the mean as C/P (see e.g. Vinkler (2010), also for a systematic approach to bibliometric measurement). Correspondingly for field-normalized citations, we can fairly easily interpret the resulting scores as a mean relative to the world (or database) average of a comparable paper. Also the h-index can be interpreted, albeit somewhat less clear and more arbitrarily (Waltman & Van Eck, 2012), as the number h of a ranked set, C, consisting of the citation counts per paper, in which h is the highest rank equal to or higher than the corresponding citation score (see also Eq. 4). For the Euclidean index, this interpretation is not the case, as the resulting indicator is on a scale with only one fixed point (0), and no intrinsic value of the unit. As an example, consider an author with two papers, of which paper a is cited ten times and paper b is cited just once. The sum of squares (SS) becomes $10^2 + 1^2 = 101$, and the Euclidean index is 10.05, but we have no means of interpreting this particular number without further information. The combination of SS and the square root hereof detaches the measurement from the original unit, C, and there is no replacement unit. In this, the Euclidean index somewhat resembles the h-index, which suffers from epistemologically parallel problems (Waltman & Van Eck, 2012).

Additionally, if paper a is cited once more, the SS becomes $(10+1)^2+1^2 = 122$, while if b is cited instead of a, $SS = 10^2+(1+1)^2 = 104$. While one might argue that this is a desirable effect and puts an emphasis on highly cited papers, it is also well-known that there is a Matthew effect leading to an overemphasis of papers that are already highly cited (Small, 2004), e.g. through perfunctory referencing. It should thus not be necessary to artificially enhance the effect of proportionally few highly cited papers.

The remaining value of the index thus becomes a ranking of authors, which is of course considered necessary by some funders and employers, but as argued before, not a recommendable practice, due to the potential implications and built-in arbitrariness of any (micro-level) citation-based ranking.

3. Empirical methodology

In this part of the paper, we will first describe the basic, bibliometric indicators to which we compare the Euclidean index, followed by a description of the robustness test and finally the dataset used for the empirical analysis.

3.1. Assessed indicators

With the aim of testing whether the Euclidean index provides new or more robust information we test its correlations with the two most basic indicators of size and impact; number of papers, P, and total number of citations, C. These two main dimensions have previously been used to test new indicators for their contributions (Bollen et al., 2009; Leydesdorff, 2009). In addition, we test correlation with the parametric combination of the two, the mean number of citations per paper. In addition to these basic indicators, we also compare the Euclidean Index to the h-index (Hirsch, 2005), R-index (Jin et al., 2007) and e-index(Zhang, 2009).

We define the six aforementioned indicators as follows:

$$P = \sum_{i=1}^{N} p_i \tag{1}$$

, is the sum of papers i for an author. Correspondingly, c_i are the citations received by paper i, which gives us the formal definition of total citations, C, for a given author (Eq. 2) and the mean of the two (Eq. 3):

$$C = \sum_{i=1}^{N} c_i \tag{2}$$

$$MC = \frac{C}{P} \tag{3}$$

By introducing a ranking function $f(C)$ which ranks the distribution of C according to c_i in descending order, we can define the h-index as the maximal index, i, which at least corresponds to c_i:

$$h(C) = \max_i \min(f(C), i) \tag{4}$$

Several of the bibliometric indicators derived from the h-index utilize what is referred to as the h-core, namely the papers ranking up to h in the above distribution. Also the R- and e-index focus on this core of publications, the former as the square root of the sum of citations (Eq. 5), and the latter by subtracting the least possible citations for h, from this same distribution, thus focusing on the excess citations (hence the name e, Eq. 6). By definition, the h-core will always contain h publications, which will be cited at least h times each. We therefore know that the h-core contains at least h^2 total citations, which can be seen in Eq. 6 as the main difference from the R-index.

$$R = \sqrt{\sum_{i=1}^{h} c_i} \tag{5}$$

$$e = \sqrt{\left(\sum_{i=1}^{h} c_i\right) - h^2} \tag{6}$$

For good measure and comparison, we also document the Euclidean index here (Eq. 7), and as can be seen from the definition, we can expect a correlation with the total sum of citations, as c_i is the independent variable in both definitions.

$$\iota_E = \sqrt{\sum_{i=1}^{N} c_i^2} \tag{7}$$

It is obvious from formulations eqs. (5) to (7) that there is an algebraic likeness of the R-, e- and ι_E-indices. The main difference is the exclusion of papers cited less than h times in the R- and e-indices and the squaring of citation counts in the ι_E-index. We should thus expect a power function correlation, with some variance due to differences between the h-core and the full publication set.

3.2. Field normalization

Normalizing citation scores to field averages can be a necessary action in order to fairly compare publications from differing fields, as the citation intensity can vary considerably from field to field. Several approaches to such normalization procedures exist, but one of the standard approaches is the Normalized Citation Score, NCS, in which all raw citation counts are divided by the average citation rate of all papers of the same type (article/review/letter), published in the same year and in the same field (Waltman et al., 2012). From this, it is possible to calculate the Mean Normalized Citation Score, $MNCS$, as the average of all NCS for a given unit of analysis. These scores are centered around 1, which is interpreted as the average performance for similar papers in the database, or simply database average. As the distribution of these scores is different from the raw citation distribution (see Fig. 1), we will also include an analysis of the correlation between field normalized Euclidean index scores and the most highly correlated other indicator.

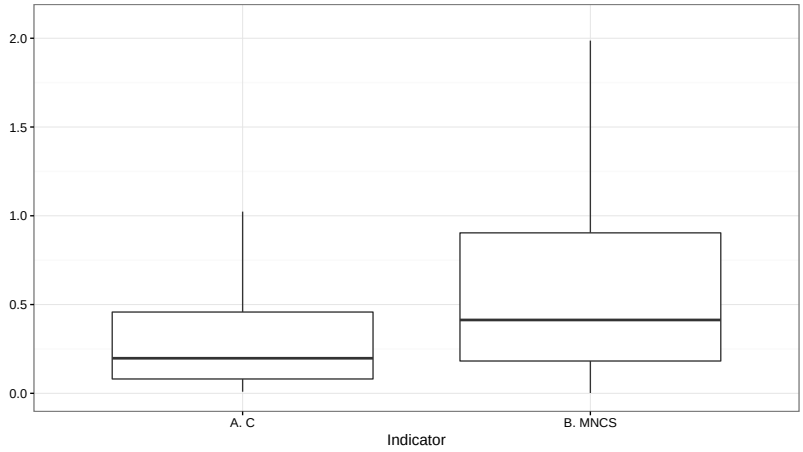

Figure 1: Frequency distributions for raw citations, C, and field normalized citations, $MNCS$. Both indicators have been rescaled to max = 100 for comparability. Rescaling was also tested using the mean of each distribution, resulting in practically similar results (not shown here).

3.3. Robustness analysis

Robustness, or stability, is an essential element of measurement. Robust indicators should provide the same value when performing a new measurement, and this interpretation makes sense in the experimental sciences from where the term originates. In fields such as bibliometrics where the sample usually is the (practical) population, repeating the experiment will always produce the same result. Bootstrapping can be a useful technique in this case, to create resamples based on the original sample/population, but by allowing for replacement, thus randomly omitting observations and repeating others for each resampling (Efron & Tibshirani, 1993). This also allows us to calculate the limits of the 95 percent interval, corresponding to confidence intervals for actual sampling experiments. In bibliometrics, there seems to be a consensus on referring to these intervals as stability intervals, as they indicate how stable the reported indicator scores are, or how sensitive to outliers in the data (Schneider & Van Leeuwen, 2014). The technique is problematic, as a) outliers in bibliometric data have a significant value, and should not be omitted and b) the population is the population[1], which defeats one major purpose of the resampling. Nonetheless, bootstrapping stability intervals provide a useful heuristic for assessing the sensitivity of a given observation. In combination with the correlation analyses, we can test both for novelty and increased robustness.

Bootstrapping was performed using the `boot` package (Canty & Ripley, 2015; Davison & Hinkley, 1997) for `R version 3.2.3` (R Core Team, 2015). `R` is also used for all other calculations and illustrations (using `ggplot2` (Wickham, 2009)).

[1]In this as previous mentions of "population", it should be interpreted as the practical population defined by the data available through the provider, in this case Thomson Reuters. This will of course differ from the real population, which we can never know.

3.4. Citation data

In order to provide realistic, empirical assessments of the Euclidean index, we use data on citation counts based on the Web of Science (WoS) Science Citation Index-Enhanced (SCI-E) and Social Science Citation Index (SSCI). However, we use a relational database version of these citation indices managed by the Centre for Science and Technology Studies at Leiden University (CWTS). Among other supplementary information, the database contains tables of author groupings, which are algorithmically generated links between unique authors and their papers (Caron & van Eck, 2014). Like any approach to constructing complete author bibliographies, the method behind this author disambiguation is not complete, however, Caron & van Eck report 95 percent precision and 90 percent recall per author, which is entirely sufficient for the present analysis. In addition, the algorithm generating these author profiles is conservative, so that new, potentially duplicate author profiles are created (even though they may be singletons) rather than assigning uncertain papers to existing authors. We therefore select only those authors who have written at least twenty papers. Subsequently we limit the publication set related to the authors to journal articles and reviews only, effectively removing letters, conference contributions etc. As a result, some authors will have fewer than twenty papers after this delineation. We further limit the dataset to only those authors with their first publication between 2000 and 2005. The intention of using these limitations is entirely to obtain a manageable set of authors, in this case resulting in $255,755$ observations. The publication years are chosen to reflect a current sample of publishing authors, while still allowing for some time to accrue citations. The selected authors have in total written $2,667,186$ papers until week 13, 2016. For all of these papers we gathered the total amount of citations up until the most recent update in 2016.

4. Results

In the following, we will first provide an overview of the relationship between the Euclidean index and the four other indicators, followed by a bootstrapping analysis, in combination to test for novelty. This is followed by a thought experiment and a simulation related to the ordinality and interpretation of the Euclidean index.

4.1. Novelty and robustness

In Fig. 2 we plot the Euclidean index as a function of the six other indicators, and the two field-normalized versions. The most clear correlation for the basic units, as anticipated, is with the sum of citations per author, C. The total papers, P, at first glance are reciprocally correlated with ι_E, but plotting ι_E against P^{-1} does not provide a clear image of a linear relationship (plot not shown). Mean citations, MC, are not correlated with the Euclidean index, but interestingly, the function creates a lower and higher threshold for ι_E depending on MC, as is also the case for C. As also documented by Perry & Reny (2016), there is no correlation with the h-index. These relationships make a lot of sense, when considering the variables involved in each indicator, and the previous findings by Leydesdorff (2009) and Bollen et al. (2009). The e- and R-indices are very clearly correlated with the Euclidean index, while it is more unclear for the two normalized indicators.

A linear regression of ι_E as a function of C gives a good estimate ($R^2 = .898$), with $\beta = .705$ but a negative intercept ($\alpha = -1.193$). The field normalized Euclidean index also correlates with the total field-normalized citation score ($R^2 = .873$, $\beta = .598$). We find the observations may reveal a power function correlation rather than a purely linear (also with regard to the design of the ι_E indicator), and therefore plot log-log transformed observations in Fig. 3, revealing a more clear correlation. We observe the same for the e- and R-indices, which are methodically even closer related, and observe very clear correlations with ι_E in both Fig. 2 and Fig. 3. Also a linear regression on the log-log

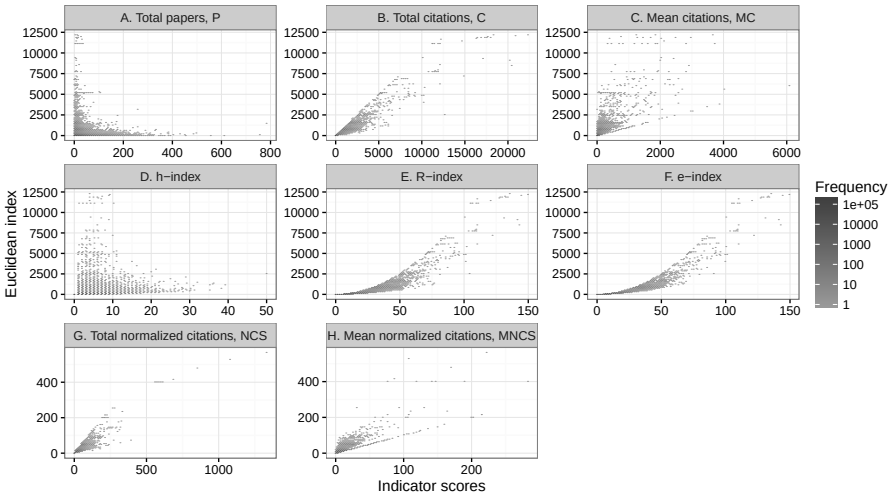

Figure 2: Correlation of ι_E with P, C, MC, h, R, e, NCS and $MNCS$.

transformed data give very strong correlations with all four indicators (see table 1), although the distribution of the sum of field-normalized citations does not seem to be as clearly correlated as the statistics indicate. This informs us that the Euclidean index does not provide us with new information, but rather transforms existing information. And while we have a rough idea about how to interpret a total citation count (although this is problematic as well), the Euclidean index is ordinal, and thus less informative, and the same can be said for the e- and R-indices. The question therefore is, whether the indicator is more robust than the total citation counts. We therefore bootstrap a subset of the authors, restricting it to those with more than 50 papers and at least one citation. This sample has $7,122$ observations, and for each observation resampling is performed with $1,000$ replications. We perform the bootstrapping on the Euclidean index, total citations and R-index, as these are the two indicators most strongly correlated with the Euclidean index. In order to appropriately compare the indicators, which operate on different scales, we rescale them to an index $100 = \max(\text{indicator})$, and the same procedure is performed on the stability intervals of each bootstrapped paper set. We use 95 percent stability

Table 1: Statistics for linear regression models, using ι_E as independent variable. For NCS, the Euclidean index is also calculated on field-normalized citations.

Indicator	R^2	α	β
Total citations $(C+1)$	.986	-.004	.921
R-index + 1	.977	-.284	1.86
e-index + 1	.976	.053	1.79
Total normalized citations $(NCS+1)$	.978	.002	.861

intervals, and report the range of the interval in Fig. 4, from which it is clear that the sum of citations is just as stable an indicator as the Euclidean index, and the R-index is less stable. Linear regression strongly estimates the sum of citations and Euclidean index to have practically identical stability intervals $(R^2 = .956, \beta = .999)$. The Euclidean index does thus not provide more robust information than the total citations, but clearly not less either. The index is clearly more stable than the R-index, especially in the lower range of stability intervals.

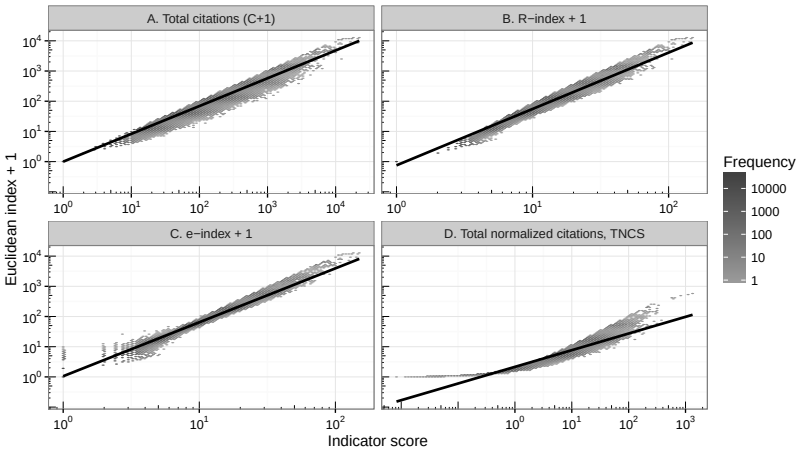

Figure 3: Correlation of ι_E and C, with regression. All indicators have been modified by adding 1, to avoid $log(0)$. Regressions are illustrated as solid, black lines).

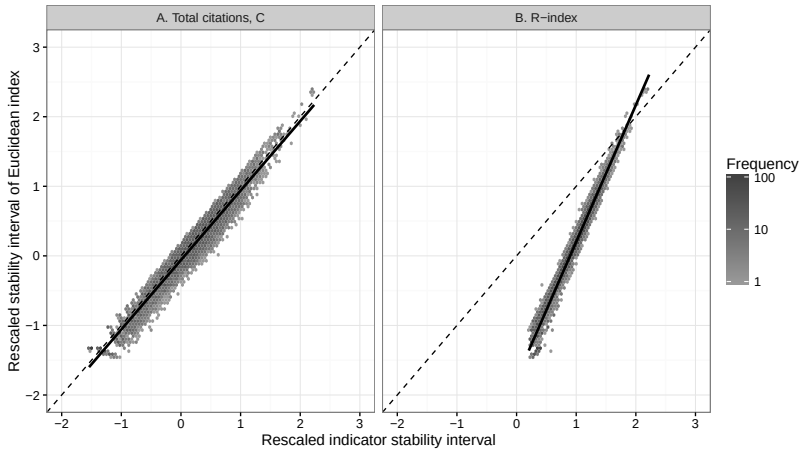

Figure 4: Bootstrap stability intervals for ι_E and C. All observations have been transformed to 100 = max(indicator) and subsequently log-transformed. Observations show the stability interval range of ι_E as a function of same for C and R. The data is restricted to those authors with at least 50 publications and at least one citation. The dashed lines represents $y = x$, while the solid lines represent regressions on the correlations.

4.2. Distinction and ordinality

With regard to the interpretation of the Euclidean index as an ordinal index, which despite its ordinality can be used to determine to which degree an author's production is a multiple of another authors' (Perry & Reny, 2016, 2728-2729), we find it useful to illustrate with a thought experiment how this is strictly dependant on the definition and context of the assessment situation. We wish to investigate two perspectives of this: what is the potential difference in ι_E for two authors with equal total citations, and what is the potential difference in total citations for two authors with equal ι_E. For the first perspective, assume an author with two papers $= c_1 = 100, c_2 = 0$ and another author with papers a and b, where $c_a + c_b = 100$. We can then make a function of c_b describing the ι_E of the second author:

$$\iota_E(c_b) = \sqrt{(100 - c_b)^2 + c_b^2}$$

It is obvious that this is a polynomial function, despite the damping effect of the square root, and the accompanying distribution is shown in Fig. 5.

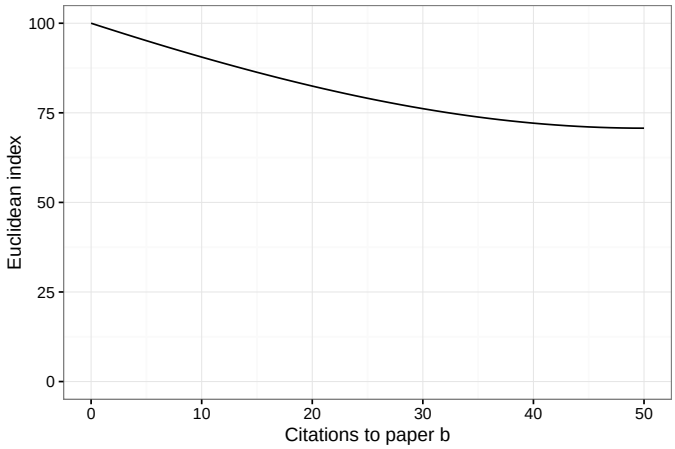

Figure 5: Polynomial function for a two-paper model.

For the second perspective, we keep a constant $\iota_E = 100$ and using the same model before of an author with just two papers, we show here the function and distribution of this situation. We can describe the citations to c_b as a function of c_a and ι_E as

$$c_b = \sqrt{\iota_E^2 - c_a^2}$$

and subsequently a function

$$f(c_a) = c_a + \sqrt{\iota_E^2 + c_a^2} = c_a + \sqrt{100^2 + c_a^2}$$

showing the sum of citations to a and b. The distribution of this function is shown in Fig. 6

From these two thought experiments we can see that two authors with the same total citations can have large differences in ι_E, with the largest differences found between authors who receive a steady amount of citations on each paper versus those with few exceptionally highly cited papers, a distinction also used in a previous analysis of the h-index, and referred to as "big producers" and "selective researchers" (Costas & Bordons, 2008). We also see, that identical Euclidean indices for two authors can have quite different distributions behind them, and that statements such as *"list x is as good as the list that has as*

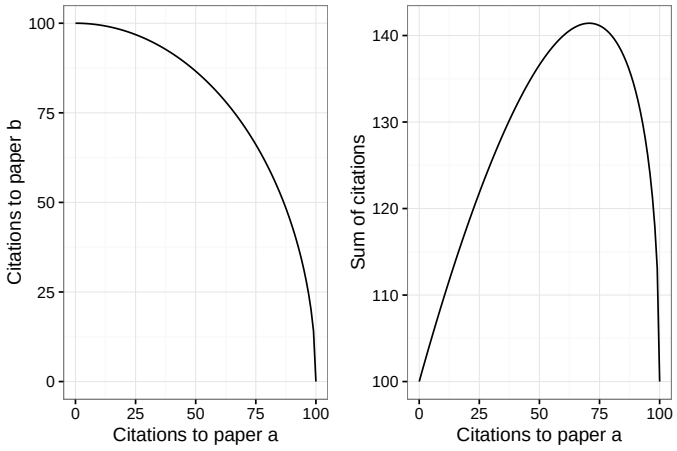

Figure 6: Polynomial function for a two-paper model.

many papers as list y, but receives λ times as many citations on each of them." (Perry & Reny, 2016, 2729) should be interpreted with extreme care, as the definition of "good" is entirely dependent on the assessment context. In a final thought experiment, we test further the degree to which the Euclidean index can be used to distinguish between two authors. By manipulating the actual citation distributions of our empirical set of authors, to all receive in total one million citations (but with the same relative distribution per author) we can compare a large set of realistically distributed authors with regard to the resulting Euclidean index. As such an analysis would be extremely sensitive to authors with just one cited paper, we limit our dataset to those with $P >= 20$. A histogram showing the frequency of index scores is shown in Fig. 7. We can see that a large number of authors are placed in the top bin, which are those with extremely narrow sets of cited papers, despite the limitations. We can safely ignore those for the case of this analysis, and observe an almost Gaussian distribution of authors, illustrating how a fixed number of citations can lead to even more different indicator scores than illustrated by the above example. The underlying effect of this is the very design of the indicator, in which citations to highly cited papers are weighted much higher.

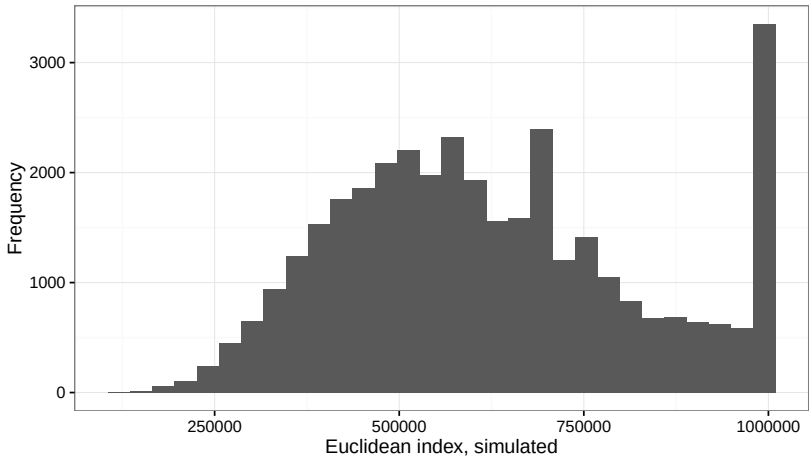

Figure 7: Authors normalised to a mega-citation. Histogram showing the distribution of Euclidean index scores for all authors, if their entire set of papers was rescaled to receive one million citations in total. Limited to authors with $P >= 20$.

5. Discussion

We have discussed and empirically analyzed the Euclidean index with respect to novelty, robustness, ability to distinguish between authors, ordinality and the interpretation of the index's scale. Our results show that the index is strongly correlated with the total citations for an author, and with similar stability intervals. The index thus does not contribute with new or more robust information, but a transformation of existing information. We can thus not recommend the use of this indicator over other, existing indicators. We have shown that authors with highly skewed citation distributions achieve higher scores on the Euclidean index than those with the same total citation count but less skewness. It is uncertain if this particular property could be applied in specific instances, but there is a clear risk of adverse effects if proper care is not taken, and we would ultimately recommend the use of a combined set of indicators, e.g. the $MNCS$ and the proportion of papers ranking in a given top percentile compared to the database, as these indicators are more informative, offering scale measurements rather than ranking schemes.

Through a thought experiment we have shown how the same Euclidean index score can be based on different distributions, and also the opposite is true. Even with extremely small samples such as the ones applied in the thought experiment, this can lead to a confusion as to the meaning of the score, and with more papers, the complexity increases. We have also shown that the index is dependent on which paper in a distribution receives a new citation, resulting in a situation in which some citations are more worth than other. Additionally, those papers who receive the greatest award from new citations are those which are highest cited already, thus reinforcing the Matthew effect already present in scholarly communication (Merton, 1968; Small, 2004).

Finally we have argued that the scale of the Euclidean index is arbitrary and thus of limited use to those it seems intended for; the research managers, funding organizations and scientists themselves. The main function of the index is as a ranking mechanism. While this also seems to be the primary intent of Perry & Reny, they also claim a scalable kind of ordinality. While this is to some extent true, it is highly dependent on the assessment situation. This is of course true for all indicators of scholarly impact, and as argued above, assessment should never rely on (one) bibliometric indicator(s) alone, especially in the case of micro-level analysis (Aksnes & Taxt, 2004; Hicks et al., 2015; Leydesdorff et al., 2016; Vinkler, 2007; Waltman et al., 2013; Costas et al., 2010).

There are several useful and sensible other approaches to micro-level bibliometric analyses, and there are indicators which can provide supporting information to peer-reviewing of individuals (Aksnes & Taxt, 2004), but as has been stated before, and as the results of our study as well as that of Perry & Reny shows, ranking individuals by any indicator for assessment leads to arbitrary selections based on purely statistical attributes of a biased indicator[2] based on imperfect data. While we support designing indicators not just for bibliometricians, but also managers and scientists etc., we also strongly oppose the concept of ranking individuals and indicators should be designed for more

[2] All indicators distill information, and through this bias certain elements.

than ordinality.

6. Acknowledgments

The author wishes to thank Jesper W. Schneider and Carter W. Bloch for helpful comments and discussion.

References

References

Aksnes, D. W., & Taxt, R. E. (2004). Peer reviews and bibliometric indicators: a comparative study at a Norwegian university. *Research Evaluation*, *13*, 33–41. doi:10.3152/147154404781776563.

Bollen, J., Van de Sompel, H., Hagberg, A., & Chute, R. (2009). A principal component analysis of 39 scientific impact measures. *PloS one*, *4*, e6022. doi:10.1371/journal.pone.0006022.

Bornmann, L., Mutz, R., & Daniel, H.-D. (2008). Are there better indices for evaluation purposes than the h index? A comparison of nine different variants of the h index using data from biomedicine. *Journal of the American Society for Information Science and Technology*, *59*, 830–837.

Braun, T., Glänzel, W., & Schubert, A. (2006). A Hirsch-type index for journals. *Scientometrics*, *69*, 169–173. doi:10.1007/s11192-006-0147-4.

Canty, A., & Ripley, B. (2015). *boot: Bootstrap R (S-Plus) Functions*. R package version 1.3-15.

Caron, E., & van Eck, N. J. (2014). Large scale author name disambiguation using rule-based scoring and clustering. In *Proceedings of the Science and Technology Indicators conference 2014* (pp. 79–86). Leiden, Netherlands: CWTS, Leiden University.

Costas, R., & Bordons, M. (2008). Is g-index better than h-index? An exploratory study at the individual level. *Scientometrics*, *77*, 267–288. doi:10.1007/s11192-007-1997-0.

Costas, R., van Leeuwen, T. N., & Bordons, M. (2010). A bibliometric classificatory approach for the study and assessment of research performance at the individual level: The effects of age on productivity and impact. *Journal of the American Society for Information Science and Technology*, *14*, n/a–n/a. doi:10.1002/asi.21348. arXiv:0803.1716.

Davison, A. C., & Hinkley, D. V. (1997). *Bootstrap Methods and Their Applications*. Cambridge: Cambridge University Press.

Efron, B., & Tibshirani, R. J. (1993). *An introduction to the Bootstrap*. New York: Chapman & Hall.

Glänzel, W., & Moed, H. F. (2012). Opinion paper: thoughts and facts on bibliometric indicators. *Scientometrics*, . doi:10.1007/s11192-012-0898-z.

Hand, D. J. (2004). *Measurement theory and practice: The world through quantification*. London, UK: Arnold.

Hicks, D., Wouters, P. F., Waltman, L., de Rijcke, S., & Rafols, I. (2015). Bibliometrics: The Leiden Manifesto for research metrics. *Nature*, *520*, 429–431. doi:10.1038/520429a.

Hirsch, J. E. (2005). An index to quantify an individual's scientific research output. *Proceedings of the National Academy of Sciences of the United States of America*, *102*, 16569–72. doi:10.1073/pnas.0507655102.

Jin, B., Liang, L., Rousseau, R., & Egghe, L. (2007). The R- and AR-indices: Complementing the h-index. *Chinese Science Bulletin*, *52*, 855–863. doi:10.1007/s11434-007-0145-9.

Leydesdorff, L. (2009). How are new citation-based journal indicators adding to the bibliometric toolbox? *Journal of the American Society for Information Science and Technology*, *60*, 1327–1336. doi:10.1002/asi.

Leydesdorff, L., Wouters, P. F., & Bornmann, L. (2016). Professional and Citizen Bibliometrics: Complementarities and ambivalences in the development and use of indicators, . (pp. 1–40). arXiv:1609.04793.

Merton, R. K. (1968). The Matthew Effect in science. *Science*, *159*, 56–63.

Moed, H. F. (2009). New developments in the use of citation analysis in research evaluation. *Archivum immunologiae et therapiae experimentalis*, *57*, 13–8. doi:10.1007/s00005-009-0001-5.

Perry, M., & Reny, P. J. (2016). How To Count Citations If You Must. *American Economic Review*, *106*, 2722–2741. doi:10.1257/aer.20140850.

R Core Team (2015). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing Vienna, Austria.

de Rijcke, S., & Rushforth, A. (2015). To intervene or not to intervene; is that the question? On the role of scientometrics in research evaluation. *Journal of the Association for Information Science and Technology*, *66*, 1954–1958. doi:10.1002/asi.23382. arXiv:0803.1716.

Rushforth, A., & de Rijcke, S. (2015). Accounting for Impact? The Journal Impact Factor and the Making of Biomedical Research in the Netherlands. *Minerva*, *53*, 117–139. URL: http://link.springer.com/10.1007/s11024-015-9274-5. doi:10.1007/s11024-015-9274-5.

Schneider, J. W., & Van Leeuwen, T. N. (2014). Analysing robustness and uncertainty levels of bibliometric performance statistics supporting science policy. A case study evaluating Danish postdoctoral funding. *Research Evaluation*, *23*, 285–297. doi:10.1093/reseval/rvu016.

Small, H. G. (2004). On the shoulders of Robert Merton: towards a normative theory of citation. *Scientometrics*, *60*, 71–79. doi:10.1023/B:SCIE.0000027310.68393.bc.

Van Raan, A. F. J. (1998). In matters of quantitative studies of science the fault of theorists is offering too little and asking too much. *Scientometrics*, *43*, 129–139.

Van Raan, A. F. J. (2006). Comparison of the Hirsch-index with standard bibliometric indicators and with peer judgment for 147 chemistry research groups. *Scientometrics*, *67*, 491–502. doi:10.1556/Scient.67.2006.3.10. arXiv:0511206.

Vinkler, P. (2007). Eminence of scientists in the light of the h-index and other scientometric indicators. *Journal of Information Science*, *33*, 481–491. doi:10.1177/0165551506072165.

Vinkler, P. (2010). *The evaluation of research by scientometric indicators*. Oxford, UK: Chandos Publishing. doi:10.1002/asi.21417.

Waltman, L., Calero-Medina, C., Kosten, J., Noyons, E. C. M., Tijssen, R. J. W., Van Eck, N. J., Van Leeuwen, T. N., Van Raan, A. F. J., Visser, M. S., & Wouters, P. F. (2012). The Leiden ranking 2011/2012: Data collection, indicators, and interpretation. *Journal of the American Society for Information Science and Technology*, *63*, 2419–2432. doi:10.1002/asi.22708.

Waltman, L., & Van Eck, N. J. (2012). The inconsistency of the h-index. *Journal of the American Society for Information Science and Technology*, *63*, 406–415. doi:10.1002/asi.21678. arXiv:1108.3901.

Waltman, L., Van Eck, N. J., & Wouters, P. F. (2013). Counting publications and citations: Is more always better? *Journal of Informetrics*, *7*, 635–641. doi:10.1016/j.joi.2013.04.001.

Wickham, H. (2009). *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York.

Zhang, C. T. (2009). The e-Index, Complementing the h-Index for Excess Citations. *PLoS ONE*, *4*, 2–5. doi:10.1371/journal.pone.0005429.